

1 Large Scale Networking (LSN) Coordinating Group Charge to MAGIC

For FY2012, the LSN tasked MAGIC to:

1. Track and promote best practices for cloud, distributed, and Grid computing
2. Track and promote best practices for Identity Management

This document represents MAGIC's response to the second task regarding Identity Management.

2 Background: Identity Management for the Open Science Community

Collaboration is at the very center of the research and education mission, with inter-institutional and international engagements a very common characteristic. Identity management is fundamental for establishing trust in these collaborations by managing entities and their privileges – who they are, how they are identified, how they are authenticated, what privileges they have, what roles and responsibilities they have – and enabling the communication of that identity information, allowing for authentication and authorizing. As collaborations have increased in scale, complexity and in their span of organizations, countries and continents, identity management is expanding to support these new collaborations.

Federated identity refers to identity management spanning multiple organizations. It allows identity information to be shared between a pair of organizations and used by entities in those organizations to establish trust in each other. This distribution of IdM across organizations leads to the need for shared policy defining the meaning of the identity information, agreed-to technologies for its communication, the acceptable use of sharing information, the degree to which it can be trusted, etc. *Federations* (e.g., InCommon) are formed amongst organizations that agree to a set of technologies and policies, hence paving the way for the trust between entities in their organizations.

Building on federated identity, collaborations need to define their membership, and the roles of that membership. They also often create their own services to serve their membership, for example web portals or compute job schedulers. This brings about the need for:

1. A set of collaboration applications, such as wikis, email lists, calendaring (ad hoc and event), and conferencing, etc., utilizing open interfaces (API's, SAML assertions, LDAP, etc) for their identity and access needs so they can be readily integrated into the collaboration ecosystem.
2. A set of domain tools (computation systems, storage, common scientific analytic tools, databases and data nets, etc.) that also have their identity and access control needs met with open community-standard interfaces for ready integration. These tools are noteworthy not only for their users, relatively small in number but consequential in importance, but their requirements: command line applications, process authentication, delegation of authorization needs, etc.
3. A collaboration-centric identity management and access control mechanism that leverages the growing research and education federation infrastructure, giving the collaborators easy tools to mingle institutional and collaboration attributes and permissions in order to manage access to collaboration resources, and provides consistent user experiences.

3 Findings

Based on its research, MAGIC identified the following findings regarding Identity Management.

- I. Identify Management has 3 parts:
 - a. Identity Providers, who certify you are who you say you are.
 - b. Attribute Providers, who certify the roles and responsibilities you have at a specific institution or inside a real or virtual organization.
 - c. Service Providers, who consume identity information from Identity and Attribute providers to serve their users.
- II. Commercial (social) identity providers have arrived on the scene, for example, Facebook and Google.
- III. Science communities have needs that are not fulfilled by the identity management ecosystem, for example:
 - a. They need command line, batch, and script-based service interfaces. While the current focus for identity management has been on web-based applications/interfaces.
 - b. Attribute providers are not emerging in the same way as identity providers, presumably due to a lack of a business model.
- IV. Applied research and development activities are needed to make technologies ready for wide scale, easy use by science communities.
- V. Technical Standards, community best practices, libraries, and API's are needed to support application developers.
- VI. Agreement is needed as to what attributes mean in the context of science VOs.
- VII. The objective is for individuals to have a single identity with multiple attributes, provided by many virtual and/or real organizations.
- VIII. International privacy guidelines need to be incorporated into deployed systems.
- IX. Assessment of compliance with US federal mandates on international collaborations needs to be considered.
- X. The risk impact on the use of federated identity management for scientific applications is not completely understood.

- XI. Previous solutions have implemented many of the identity management elements, but have not gained widespread traction outside their development communities.

4 Future Actions

Based on its findings, MAGIC members identified a need for future actions.[VW1]

- I. The basic technical strategy (separating identify providers from attribute providers) has laid the foundation for widespread implementations of Identity Management systems. However, identifying policy issues (national and international privacy guidelines, government implementation mandates, risk assessments, etc.) or best practices for operating these IdM systems remain topics for discussion.
- II. Web-based systems, where a human is able to interactively respond to prompts or security challenges, do not meet all the needs of the scientific user community. Enhancements that address the need for non-interactive usage need to be widely integrated and deployed into IdM systems.
- III. Applied research and development activities that focus on wide scale deployment, ease of use (for users, system administrators, and application developers), and backward compatibility with legacy applications are all challenges that must be overcome to achieve wide spread acceptance of IdM systems.
- IV. Identity systems such as InCommon are making progress in the support of science communities with efforts such as the InCommon Research and Scholarship category, but room for improvement exists to make its services more readily usable by virtual organizations and the service providers that represent them.

5 Notes from Information Gathering Meetings

This section captures the notes from the relevant monthly MAGIC teleconferences that serve as the basis for the preceding Findings and Future Actions. For the complete materials, please see the MAGIC WIKI at: http://connect.nitrd.gov/magicwiki/index.php?title=Meeting_Minutes_and_Materials

5.1 August 3, 2011 MAGIC Teleconference

Updates on Internet Identity: Ken Klingenstein

Internet identity has become pervasive in two flavors:

- A rapidly growing and maturing federated identity infrastructure used extensively in the R&E sector globally
- Theoretically interoperable social identity providers serving large masses of social and low-risk applications

Federated uses vary significantly by country and sector (medical, real estate are building corporate federated identities)

Social identities have been used beginning in 2007. A small number of major players share a set of non-interoperable deployments of weak protocols. A move is being made to convergence around OpenID

Connect. An integration of federated and social approaches is emerging including Social2SAML gateways.

The national Secure Transactions in Cyberspace (NSTIC) is a White House initiative on citizen-government security/privacy. See: www.nstic.gov/nstic. It works well with SAML and R&E federations.

Federated identity is still a work-in-progress. Major issues remaining include non-Web applications, inter-federation and developing an attribute ecosystem.

InCommon is a federation of 250+ universities, 450+ participants of over 10 million users that is growing rapidly. Certificate services bind InCommon trust policies to new applications.

Important new InCommon services include: Research.gov, electronic grants administration from NIH, CILogon, IEEE, Educause, NBCLearn, University Tickets, and many others.

New developments for InCommon include:

- Growth and managing service
- Silver and higher levels of assurance/service
- uApprove end user control of attribute management
- Social2SAML coordination
- Personal certificates

5.2 October 5, 2011 MAGIC Teleconference

Identity Management for Distributed Science,” Von Welch, Center for Applied Cybersecurity Research, Indiana University

- Progress is being made
- Adoption is slow. Change is slower after initial deployment
- Recommendations for MAGIC
 - Foster international interoperability
 - Define community requirements
 - E.g. LOA comes from risk, risk comes from assets, which are increasingly data, There are no data security needs assessments
 - E.g. Should we be leveraging outside IdM rather than rolling our own? If so, what would we need from InCommon, OpenId, NSTIC, etc.?
 - Monitor Moonshot/SAML ECP and jump in and support the winner at the appropriate time?

Open Grid Forum (OGF) Status: Alan Sill

- International Grid Trust Federation (IGTF) is the right forum to coordinate trust issues not OGF
- Federated security – writing profile use of access space to include clouds; this is not a technology question,

- Motivation – We want to expand the user community/projects by orders of magnitude
- Formulate profile and hope to use existing standards, e.g., if you wish to connect via shibboleth you need to express the membership in a virtual organization
 - First work product available in grid architectures
 - Second area – recent events in CC, resistance to mod architectures, want to know how to adapt their system distributed I federation
- Automation of service agreements
- How new standards can be folded into existing standards?
- Need rapid response (small scale workshops are a good approach)
- MAGIC – craft vehicles to respond to needs

5.3 January 4, 2012 MAGIC Teleconference

Identity Management Systems for Collaborations and Virtual Organizations; Ken Klingenstein

Consumer marketplace identity management is led by Google with participation by Paypal, Yahoo and others. It is based on the Open Internet Exchange (OIX) using a new standard OpenID Connect. OpenID Connect is based on Shibboleth with additional capabilities in JSON. It uses SAML attributes and metadata that enables integration. Other major players are sitting on the sidelines including Facebook and Twitter. ISOC is interested in moving forward cooperative Identity Management internationally regardless of what is happening in the U.S. It fosters a comprehensive integration of roles and communities.

Consumers are individuals with their roles and attributes. They retain their identity even when they assume different roles, different policies and different governance.

NSTIC (<http://nist.gov/nstic/>) provides a well-crafted architecture and approach. OMB issued a fall 2011 directive that the Federal agencies should move to external identities where appropriate. An IDTrust Conference will be held in Gaithersburg March 13-14, 2012.

InCommon currently has 250+ universities and 450+ participants and continues to grow rapidly with over 10 million current users. It has 300 university providers versus 4 providers for Google. New uses are being developed for InCommon including Wikis, shared services, cloud services, calendaring, command line apps, UHC and others. Certificate services bind the InCommon trust policies to new applications including signing and encryption. FICAM certified at LOA 1 and 2 (bronze and silver). New InCommon developments include uApprove (end user attribute management), Social2SAML coordination and personal certificates for authentication, signed mail, signed documents, encryption, etc. Silver service provides a higher level of assurance to support financial and other valued resources. Silver service is used for grants administration, TeraGrid, OSG and medical records.

Basic attributes for science applications include: high-level affiliation, opaque, persistent and non-correlating identifiers (ePTID), a persistent and human-suable identifier (e.g. kjk@internet2.edu), name, email address and an open-ended set of entitlements assigned by the institution including group membership. Attributes tend to travel in bundles. For research and scholarship (R&S) the bundle contains: name, email, authenticated identity, and affiliation.

Approaches are being developed for non-Web applications. Challenges for this space include discovery, trust anchors in the clients, attribute release and privacy management. Three categories of approaches include:

- Moonshot- GSS over Radius and maybe SAML
- Oauth and OpenID Connect
- SAML ECP (extended client profile)- Kitten

There are no turnkey deployments yet.

Inter-federation provides connection among autonomous identity federations. This is critical for global scaling, accommodating state and local federations and integration across vertical sectors. Operational capabilities include Kalmar2 Union and eduGAIN. Key technologies are being developed and used: PEER, metadata enhancements and tools and discovery.

Virtual Organization Identity Management

There are three contexts for VO ID Management:

- Internet-scale
- Campus/enterprise
- Virtual Organization

Primary issues are how to leverage the Internet and enterprise to serve the VO. Including leveraging security, privacy, efficiency, ease of use, sustainability... while identifying and engineering what is unique about the VO

Collaboration Management Platforms include SurfNet and COmanage.

5.4 Notes from other Discussions

- There are 2 scenarios for managing IDManagement for VOs. For LIGO, groups from universities (MIT, Caltech) write up an MOU with other universities for access to resources. They join wholesale into the VO. Under an alternative scenario. External to the VO set up collaboration for a specific application, e.g., sharing notes, findings, postings, developing research papers. Use resources such as CoManage to establish the collaboration.
- Institutions have to be primary for identity management for specific applications, e.g., for the Higgs Boson collaboration.
- How do we get VOs working across many organizations and apps? Identity management and attribute authorities are outsourced to assert individual identities. Your identity is established close to your home institution but is adopted for other uses like VOs. Universities are looking to the future where they rely on Google identities. The Google identity is then decorated with university attributes.
- Assume that Globus will not be replaced. The movement is toward acceptance of outside entity authorities.
- Non-Web ID Management: Current development is all for Web-based apps. There are no turnkey non-Web uses yet. Discovery and attribute release need to be based on SSH. You can

bring attributes into the local service. How are the attributes retrofitted to the local app? Is there a reasonable mapping between what attribute providers can provide and what the application server can use? Federal agencies have a role to provide this bridge. A modest investment might provide a real gain. COmanage and Assert Connect can manage attribute release. MAGIC could profile existing standards

- OGF has a Federated Security Group that deals with integrating ID Management capabilities and establishing security standards.
- Open Science Grid and European Grid projects are developing a common document to identify how federated security is provided across US and European groups.